

REMARKS

Reconsideration of the rejections set forth in the Office Action is respectfully requested. By this Amendment, the specification has been amended and claims 21, 23, 26, 34-35, 39, 41-42, and 44-46 have been amended. A drawing correction authorization request is filed concurrently herewith to make a minor change to Figure 1. Currently, claims 21-46 are pending in this application.

Objection to the Drawings

The Examiner objected to the drawings because the reference character “16” in figure 1, according to the Examiner, did not clearly show the “information sets” as was described in the specification at page 9, line 14.

The information set, described for example in the specification at page 8-9, is a collection of data associated with the communication devices and links and their defined network features. (See e.g. Page 8 at lines 3-5). The data in the information set 16 may change over time as network elements are added, links are changed, or other topography changes occur in the underlying network layout 3. (See e.g. Page 10 at lines 1-5). The information set is a particular instance of the information at a point in time. Thus, over time, there may be multiple information sets generated which may be stored, etc. At any particular time, however, there will be one instance of the data which will represent the current information set. The information set may contain multiple types of information associated with the network layout. For example, the information set may contain logout data, element state, configuration data, connectivity data, categorization data, and status and performance information. (See Page 9, lines 15-17).

Fig 1 shows a system for representing a network layout. Typically, in a system diagram, data flowing between processing blocks is illustrated with an arrow, as it was in the original Fig. 1. Specifically, original Fig. 1 shows information set data (arrow 16) flowing from the network layout to the data collector 4 and global database 5. As an attempt to clarify what data is being collected, applicants propose to include a dashed line surrounding the network elements and links in the network layout 3, and to include a reference character 16 connected to that dashed line. Applicants believe that this change will illustrate to a viewer that the information contained in arrows 16 relates to the underlying network elements. To formalize this change, applicants submit herewith a Drawing Correction Authorization Request containing the proposed drawing

correction. This drawing correction does not add new matter to the application but rather clarifies that which was originally disclosed. Additionally, as discussed in greater detail below, applicants have amended the specification to clarify that the information set contains an instance of data associated with the underlying network layout 3. In doing so, applicants have attempted to avoid use of the plural form of the term “information sets”, although multiple information sets may be generated over time as mentioned above. Accordingly, applicants respectfully request that the objection to the drawings be withdrawn.

Objection to the Specification

The Examiner provided several specific objections to the specification indicating that particular changes should be made. Applicants have amended the specification in a manner consistent with that suggested by the Examiner.

Additionally, with respect to the use of reference character 16, applicants have amended the specification to use the reference character in a consistent manner. These amendments do not add new matter to the application but rather clarify that which was originally disclosed. Accordingly, applicants respectfully request that the Examiner approve entry of these amendments and withdraw the objection to the specification.

Objections to the Claims

The Examiner objected to claims 34, 39, 44, and 45 as containing several informalities. Applicants have amended the claims in a manner believed consistent with the Examiner’s helpful suggestions and respectfully request that the objection to the claims be withdrawn.

Rejection under 35 USC 112, second paragraph.

Claims 26, 41, 42, and 46 were rejected under 35 USC 112, second paragraph, as indefinite for containing insufficient antecedent basis for several of the terms. Applicants have amended claims 26, 41, 42, and 46 to provide antecedent basis for the several terms identified by the Examiner and respectfully request that the rejection of these several claims be withdrawn.

Rejection under 35 USC 101

Claims 21-34, and 44-46 were rejected under 35 USC 101 as being directed to non-statutory subject matter. Specifically, the Examiner has taken the position that the language of the claims “raises a question as to whether the claim is directed merely to an abstract idea that is not tied to a technological art, environment or machine which would result in a practical application producing a concrete, useful, and tangible result to form the basis of statutory subject matter.”

MPEP 2106 describes in detail the categories of computer software that form statutory and non-statutory subject matter. Specifically, MPEP 2106(IV)(B)(2)(b) describes the requirements for statutory process claims. All of the rejected claims (21-34 and 44-46) are process claims and, as such, are to be analyzed according as described in this section of the MPEP.

Claim 21, as originally presented, recited:

“A method for enabling differential visualization of a plurality of aspects of a telecommunication network, said method comprising the steps of:

presenting a background image representation of at least a first of the aspects of the telecommunication network; and

presenting a foreground image representation of at least a second of the aspects of the telecommunication network over the background image representation.”

This claim is statutory for at least two reasons. First, the claim recites that the method manipulates data that represents physical objects or activities. (See MPEP 2106(IV)(B)(2)(b)(i)). Specifically, claim 21 recites that the background image representation and foreground image representation are both representations (i.e. data objects) associated with aspects of the telecommunication network. The separation of the data representing the physical objects into foreground and background image representations renders the claim statutory.

Second, the method in this instance recites a process that achieves a practical application. Specifically, claim 21 recites “a method for enabling differential visualization”. Differential visualization allows a display to be created that will allow a user to distinguish the background image representation and foreground image representation. This is a practical application and not merely an abstract idea or mathematical algorithm. Accordingly, claim 21, and those claims dependent thereon, recite statutory subject matter.

To clarify that the claim is statutory in the sense that the process is being carried out by a computer, and not a human, applicants have amended claim 21 to recite “A method for enabling differential visualization on a display.” Support for the term “display” may be found, for example, at page 8, line 12. In view of this amendment, applicants respectfully submit that claim 21 and the claims that depend on claim 21 recite statutory subject matter.

The same rationale applies to the analysis of claim 44. Specifically, claim 44 recites a method for presenting a visual representation of a telecommunication network layout. The claimed method, in this instance, both manipulates data that represents physical objects or activities, and recites a practical application rather than an abstract idea or bare algorithm. Specifically, claim 44 recites the step of generating a representation of at least a portion of the information set which, as recited earlier in the claim, contains information relevant to the telecommunication network layout. Thus, the data manipulated by the claimed method represents physical objects, namely the telecommunication network layout. Additionally, the practical application in this instance includes generating a representation having background and foreground image portions. These image portions allow a user to distinguish the background from the foreground to thereby allow the operator to more easily understand the network layout. To clarify that this method is happening in a computerized environment and not in a person’s mind apart from the computer, applicants have amended claim 44 to recite “A method for presenting a visual representation on a display” For all these reasons claim 44, and those claims dependent thereon, recite statutory subject matter.

Applicants respectfully submit that the claims, as originally drafted and as amended, recite statutory subject matter. For all these reasons, applicants respectfully request that the rejection of the claims under 35 USC 101 be withdrawn.

Rejection under 35 USC 102

Claims 21-42 and 44-46 were rejected under 35 USC 102 as anticipated by Cox et al. (SIGMOD Record, Vol. 25, No. 4, December 1996). This rejection is respectfully traversed in view of the amendments to the claims and the following arguments.¹

Cox teaches a system for visualizing geographic networks on a display, by using three-dimensional graphics. There are two different systems in Cox – the first system, described in section 2.1, teaches that arcs may be drawn on a globe, while the second system described in section 2.2, teaches that arcs may be drawn on a flat map rendered in three dimensional space. The manner in which the arcs are controlled in the two systems is different given the different way of showing how the arcs interconnect points in physical space.

In the display that looks like a globe (Section 2.1) Cox teaches that one problem with using a globe to show the interconnection is that it is not possible to see the back side of the globe. (section 2.1, paragraph 3, lines 1-4) Accordingly, Cox teaches three ways to make this possible. (section 2.1, paragraph 3, lines 4-5) First, Cox teaches that it is possible to make the globe partially transparent. (section 2.1, paragraph 3, lines 5-11). By making the globe more or less transparent, it is possible to see or hide the termination points on the back (hidden) side of the globe. Thus, the translucency control in section 2.1 of Cox relates to the globe, not the arcs.

The second way to distinguish between arcs when using a Globe, according to Cox, is to allow the user to select between several different ways of routing the arc paths. (Section 2.1, paragraph 3, lines 11-17) For example, Cox allows the user to cause the routes to run through the center of the globe. Cox teaches that routing the routes through the center of the globe is useful “when used with the translucency control.” (Section 2.1, paragraph 3, lines 16-17).

The third way Cox teaches to differentiate between arcs is to cause arcs to be filtered to select only those with certain attributes. (Section 2.1, paragraph 3, lines 17-18). Cox continues in this section to state that “incorporating user interface controls such as filtering and translucency can further reduce the visual complexity of the display, and thereby lead to greater insights.” It is clear in this section that the “translucency” being referred to is the translucency of

¹ The copy of the Cox reference contains a figure (Fig. 1) that was relatively difficult to read. Applicants searched on the Internet for a better copy and located a copy of a paper by Kenneth C. Cox, et al. that is represented on its face to have appeared in Sigmod Record, Volume 24, Number 4, in December of 1996. The figures in this paper are much clearer than the version provided with the Office Action. Accordingly, applicants submit for consideration this alternative copy of the reference. Applicants will refer to the originally cited copy of Cox in response to the rejection.

the globe, not the translucency of the arcs. Indeed, Cox does not teach, in section 2.1, that the arcs on a globe may be translucent. Cox also does not teach, in section 2.1, that different arcs may be made translucent while others are rendered in saturated color. Rather, Cox either displays arcs or doesn't display arcs according to the filtering, and allows the globe to be made translucent if desired. Filtering arcs to cause the arcs to sequentially appear on the screen is not the same as showing some of the arcs in the background and others in the foreground but rather is akin to showing multiple views of the same network serially as described in the background of this application.

The second embodiment discussed in Cox is where arcs are created between points on a flat map rendered in three-dimensional space (Section 2.2). Links having the highest traffic flows are rendered to be the highest (Section 2.2, paragraph 2, lines 15-18). Cox teaches that sliders may be used to vary the properties of the arcs. (see Section 2.2, paragraph 3, line 1). In paragraph 3 of section 2.2, Cox teaches that the arcs may be rendered with translucency to avoid occlusion of otherwise hidden arc segments. (Section 2.2, paragraph 3, lines 10-12).

Cox does not teach that the user may select several of the arcs to be translucent and others to be not translucent. Rather, Cox merely states that the arcs (all the arcs) may be rendered with translucency. This is supported by the version of the paper attached hereto. Specifically, although Fig. 2 is mostly not visible in the cited version of the paper, in the attached version Fig. 7 shows a screen shot of the user interface in which a slider is provided to control the arc transparency level. Thus, it is clear from a reading of Cox that the transparency of the arcs is controlled as a whole, not on an arc-by-arc basis. Specifically, Cox does not teach that some of the arcs may be rendered transparent while other arcs are rendered in saturated color.

Cox also has a section entitled "Drill-Down Network Views" (Section 2.3) in which he teaches three different ways to display additional details about a particular node. The first view looks like spokes on a wheel, the second view looks like spokes emanating in a helix, and the third view looks like tessellates on the surface of a sphere. Better copies of the figures are contained in the attached version of the paper (See Figs. 4-6 of the attached paper). These views are completely different than the original three dimensional view of the network in which arcs showed the interconnection of the nodes on the globe.

The present invention relates to a method and system for representing and maintaining network layouts. As networks get complicated, the amount of information that may be displayed

on a management display system increases, which makes it more difficult to isolate any given aspect of the network architecture. To differentiate the different aspects of the network, network management systems generally used different views to sequentially display different aspects of the network. Thus, in previous systems, if a network manager wanted to understand how different aspects of the network worked together, the network manager would be required to sequentially view screens showing the different aspects of the network. (See e.g., Specification at page 3, lines 22-27).

The Cox system thus operates in a manner similar to the art described in the background section of this application. Specifically, as described in Section 2.1 of Cox, if the network manager would like to see different aspects of the network, the network manager may cause arcs to be filtered so that only particular arcs are displayed. In practice, as new filter rules are applied, a sequence of screen views will appear on the screen showing different aspects of the network. The drill-down views in Cox operate the same way – if an user wants to see additional information about traffic to a particular country, the user may click on the icon for that country to cause a new view in the form of a spoke view, helix view, or pincushion view, to be displayed.

Applicants discovered that it was possible to effectively present network information by causing a portion of the information representing the telecommunication network to be displayed as a background image and another portion of the information representing the telecommunication network to be displayed as a foreground image. For example, the background image may be presented using a grayed out image while the foreground image uses saturated colors. By allowing the network operator to toggle information from the background into the foreground, and vice-versa, applicants discovered that it is possible to move through multiple views while keeping the network information in context. Cox does not toggle information from the foreground to the background, but rather toggles information on and off the screen using filtering.

The Examiner has taken the position that Cox anticipates claim 21 which recites (as amended):

Claim 21. A method for enabling differential visualization on a display of a plurality of aspects of a telecommunication network, said method comprising the steps of:

presenting a background image representation of at least a first of the aspects of the telecommunication network; and

presenting a foreground image representation of at least a second of the aspects of the telecommunication network over the background image representation, said second of the aspects being user-selectable.

Specifically, the Examiner has taken the position that Cox teaches “translucent arc segments, globe, and arcs that lie spatially behind other arcs represent the background image and the arc segments that aren’t translucent or that spatially lie in front of other arcs represent the foreground image.” Applicants respectfully submit that Cox does not teach the simultaneous use of translucent and non-translucent arcs.

As discussed above, Cox teaches that a globe may be made translucent. In this view, the arcs are solid lines and the translucency of the globe is user-adjustable to allow the termination points to be viewed. In an alternative presentation, Cox teaches that a global map may be used with arcs drawn between the various cities. The arcs, in this presentation, may all be made translucent at the same time by the user. Thus, Cox does not teach one display in which the user is allowed to select some arcs to be translucent and other arcs to be rendered using a solid color.

Applicants have amended claim 1 to recite that the second aspect of the telecommunication network that is to be presented as the foreground image is user selectable. The globe, the arc segments that lie behind the globe, and the arcs that lie spatially behind the other arcs in Cox are not user selectable. Thus, this amendment to claim 1 differentiates the claimed subject matter from Cox so that Cox fails to anticipate claim 1 as amended.

Additionally, claim 21 recites that the background image contains at least a first of the aspects, and that the foreground image contains at least a second aspect. Cox teaches a way to visualize only one aspect on the display at a time, namely the amount of traffic between different parts of the network. Thus, Cox does not teach a way to enable differential visualization of at least first and second aspects of the telecommunication network as claimed in claim 21. Accordingly, applicants respectfully request that the rejection of claim 21, and dependent claims 22-34, be withdrawn

Claims 35-43

With respect to claim 35, the Examiner has taken the position that the “translucent arcs or the arcs with differing colors may represent the reference view, and the non-translucent arcs or

the arcs with a certain color the user is interested in viewing represents the overlay view of a managed telecommunication network.”

As discussed above, Cox does not allow the user to select particular arcs to be translucent and others to be non-translucent. The portion of Cox cited as support for the Examiner’s position (Section 2.1, lines 13-15, and 18-20, and Section 2.2, paragraph 3, lines 10-12) don’t actually teach a system in which the user may select particular arcs to be transparent and others to be non-transparent.

Section 2.1 teaches a visualization system in which the traffic patterns are displayed as arcs of different colors around a globe. The globe may be made translucent if desired by the user. No-where in Section 2.1 does Cox teach that the arcs may be made translucent; the only translucent aspect of the visualization in Section 2.1 is that of the globe.

Section 2.2 teaches “Another related technique” (See Section 2.2, paragraph 1, line 1), i.e. a technique different than that described in Section 2.1, in which arcs are drawn among nodes positioned on a flat 2D map rendered in a three dimensional space. The arcs (all arcs together) may be rendered with translucency, or all arcs may be rendered in saturated color. The amount of translucency may be controlled by the user, but the translucency applies to all arcs together as a whole; individual arcs may not be selected by the user to be selectively translucent or non-translucent.

Applicants have amended claim 35 to recite that the GUI is configured to provide a network manager with an ability to simultaneously display a reference view of a managed telecommunication network and an user-selectable overlay view of the managed telecommunication network in a distinguishable fashion in the window. Since Cox does not teach a GUI in which the network manager may simultaneously view both a reference view and an user-selectable overlay view, applicants respectfully request that the rejection of claim 35, and dependent claims 36-43, be withdrawn.

Claims 44-45

With respect to claims 44-45, applicants have amended claim 44 to recite that the foreground image is indicative of at least a second aspect of the telecommunication network layout, the foreground image being user-selected and derived from at least a second data subset of the information set. Since Cox fails to teach or suggest a system in which the foreground

image is user-selected, applicants respectfully request that the rejection of claims 44-46 be withdrawn.

Claim 46

Applicants note that claim 46 as originally presented recited that at least one of the first and second data subsets are user selectable. In rejecting this claim, the Examiner took the position that Cox teaches that the user can select which arcs are in the foreground and which are in the background. As described above, Cox does not teach this feature. Rather, the control in Cox is of all arcs together. This is shown more clearly in Fig. 7 of the version of the paper attached hereto, which appears to have been Fig. 1 of the figure in the cited version of the paper. As shown in Fig. 7, Cox has an arc translucency control which works as a slide to control the arc transparency. It is clear by looking at this figure that the transparency of the arcs is controlled as a whole, not on an arc-by-arc basis. Thus, applicants respectfully request the Examiner to withdraw the rejection of claim 46.

Rejection under 35 USC 103

Claim 43 was rejected under 35 USC 103 as unpatentable over Cox in view of Uribe (U.S. Patent No. 5,461,560). Claim 43 depends from independent claim 35 and, as such, is patentable for at least the reasons set forth above.

Conclusion

In view of foregoing amendments and remarks, it is respectfully submitted that the application is now in condition for allowance and an action to this effect is respectfully requested. If there are any questions or concerns regarding the amendments or these remarks, the Examiner is requested to telephone the undersigned at the telephone number listed below.

If any fees are due in connection with this filing, the Commissioner is hereby authorized to charge payment of the fees associated with this communication or credit any overpayment to Deposit Account No. 502246 (Ref: NN-14538).

Respectfully Submitted


John C. Gorecki
Registration No. 38,471

Dated: September 20, 2005

John C. Gorecki, Esq.
P.O. Box 553
Carlisle, MA 01741
Tel: (978) 371-3218
Fax: (978) 371-3219

3D Geographic Network Displays

Kenneth C.Cox (kcc@research.bell-labs.com)
Stephen G. Eick (eick@research.bell-labs.com)
Taosong He (taosong@research.bell-labs.com)

November 1996

This paper appeared in *Sigmod Record*, Volume 24, Number 4, December 1996.

Copyright © 1996 by the Association for Computing Machinery, Inc. Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Publications Dept, ACM Inc., fax +1 (212) 869-0481, or permissions@acm.org.

3D Geographic Network Displays

Kenneth C. Cox, Stephen G. Eick*, Taosong He
email: {kcc,eick,taosong}@research.bell-labs.com
Bell Laboratories

17 September 1996

Abstract

Many types of information may be represented as graphs or networks with the nodes corresponding to entities and the links to relationships between entities. Often there is geographical information associated with the network. Common examples include telecommunications networks, airline or other traffic routes, and distribution systems for electricity, water, and other resources.

The traditional way to visualize geographic networks employs node and link displays on a two-dimensional map. These displays are easily overwhelmed, and for large networks become visually cluttered and confusing. To overcome these problems we have invented five novel network views that generalize the traditional displays. Two of the views show the complete network, while the other three concentrate on a portion of a larger network defined by connectivity to a given node. Our new visual metaphors retain many of the well-known advantages of the traditional network maps, while exploiting three-dimensional graphics to address some of the fundamental problems limiting the scalability of two-dimensional displays.

*Correspondence contact: Room 1G-351, Bell Laboratories, 1000 East Warrenville Road, Naperville, IL 60566,

1 Introduction

With the explosive growth of networks and widespread availability of the World Wide Web, huge volumes of information are now accessible in networked databases. This has occurred because of ongoing technology trends such as the development of larger, cheaper disks, the widespread availability of mass storage systems, and the overall computerization of business and government. Now that people are gaining widespread access to data warehouses, another problem is emerging: how to extract the information latent in the data. Unfortunately, our ability to extract knowledge from large databases has not kept pace with the increasing data volumes.

One approach for extracting useful information from the data stores involves *visualization*. Visualization means presenting the complex datasets in pictorial forms through interactive graphics, and using human recognition capabilities to extract meaningful information from them. Humans have evolved a sophisticated visual system and excel at mining information and patterns from complex images. Information Visualization exploits this capability by encoding data using color, shape, position, texture, motion, etc., and rendering the result on a graphics workstation. Visualization links the two most powerful information processing systems—the human mind and the modern computer.

The fundamental problem in Information Visualization involves inventing the visual metaphors for representing the non-physical data. The visualization of data about physical phenomena, whether obtained from simulations or from sensors, has an extensive literature (e.g., [PvW93] and [Kau91]) and well-developed techniques, many of which are available as commercial software. The construction of these visualizations is greatly facilitated by the fact that they represent physical entities, meaning that a natural representation of the data is available in almost all cases. Abstract information, however, usually has no natural form or representation, and the literature and techniques are correspondingly less well-developed. The visual *metaphors*, that is, the representations for making abstract relationships visible, are only now emerging [Tuf83] [Tuf90] [GE95] [GE96].

Graphs are recognized as an important visual metaphor for many different types of information [RCM93] [SB94] [LR94] [Noi93]. Graphs consists of nodes and links; there are often statistics associated with the nodes and links. Many types of information may be represented by graphs with the

nodes corresponding to entities and the links corresponding to relationships among the entities. The properties of the graphical objects used to draw the graph, such as the color and size of glyphs or the thickness of lines, can be used to encode statistics about the nodes and links. The interesting aspects of a graph often involve its topology, structure, and connectivity, and the positioning or layout of the nodes and links can often be exploited to emphasize these properties.

Positioning nodes and links for drawing arbitrary graphs is a delicate operation involving subtle trade-offs [GKNV93] [BETT94]. There is a fundamental perceptual issue involving the node placement and the length of the links connecting them. Nodes that are close to each other are perceived to be related and yet the link between them does not receive much visual real estate, making it hard to see. Conversely, distant nodes are not perceived to be related and yet are connected by a long, visually-apparent link. Some of the best graph-drawing algorithms use this perceptual tension for conveying information related to particular applications.

Our interest and motivation for visualizing graphs comes from analyzing many networks associated with Bell Laboratories. In our usage, a network is a graph where the associated statistics represent *traffic*. The traffic may represent a time-varying statistic, as with the call flow in a telecommunications system throughout the day; a constant, as with network capacity; or a stochastic statistic, as with the number of IP packets sent between routers on a backbone data network.

Networks often have a geographic component, and some of the properties of the networks cannot be understood without reference to this component. The most common technique for visualizing geographic networks involves node and link diagrams [Ber81]. Glyphs (graphical objects) representing the nodes are positioned on a geographic map, with lines drawn between them showing the connections between nodes. The color and thickness of the lines may be used to represent the traffic, with the thicker and brighter lines showing the links carrying the most traffic, with the greatest capacity, and so forth. The glyphs may be colored, shaped, and sized to encode statistics associated with the nodes, for example, the router capacity, utilization, and packet losses.

Node and link diagrams are effective for visualizing small sparse networks.¹

¹A small sparse network may contain ten to a hundred nodes.

For larger and richer networks, however, the displays are easily overwhelmed and become visually cluttered with too many line crossings. One approach to solving this problem involves graph layout algorithms. The essential idea is to position the nodes spatially in ways that reduce the number of line crossings, drawing the links using carefully routed curves, and using other techniques to minimize display clutter.

These techniques are effective for showing graph structure and produce pleasing visual displays. However, they are often not suitable for geographic networks, where the position of the nodes cannot be varied without losing the geographic context. Our focus, then, is on layouts that preserve geographic context.

2 3D Network Displays

One way to solve the clutter problem inherent in 2D network displays is to draw the network in 3D. The idea is that by positioning the nodes in 3D and drawing arcs instead of links, we can eliminate the line crossings that confuse 2D displays [SA94]. Of course when viewed from any particular angle certain links may appear to cross on a 2D computer screen. The advantage, however, of the 3D representation is that through our preattentive depth perception we will automatically perceive the display correctly.

There has been previous research into 3D network displays; for some of the early papers see [FPF88] [CT96]. We build on this work and differentiate our results in three important ways:

Geographic Context. We concentrate on 3D networks layouts that maintain geographic context. As we indicated above, for many information spaces involving spatial locations, maintaining the geography in the visualization aids in the understanding.

Drill-Down Network Views. We present a suite of novel *drill-down* views that show details on demand for particular user-designated nodes.

Restricted Navigation. In general navigating a 3D network display is difficult and a user can easily become lost and disoriented. In our views we have carefully restricted the user's ability to navigate and thereby

prevent disorientation, while simultaneously striving to maintain the perceptual advantages of 3D layouts.

2.1 Global Networks

The first technique, initially described in [CE95], positions the nodes geographically on a globe and draws lines or arcs among them (see Figure 1). This results in a pleasing and informative display, looking somewhat like international airline routes, that retains the spatial information associated with the nodes and also eliminates the line crossings associated with 2D displays.

Figure 1 shows one frame from an animation of Internet traffic between fifty countries over the NFSNET/ANSnet backbone for one two-hour period during the week of February 1-7, 1993. The dataset contains the packet counts, by two-hour period, between each pair of countries. Each country is represented by a box-shaped glyph that is scaled and colored to encode the total packet count for all links emanating from the country. The glyphs are positioned at the locations of the countries' capitals and extend perpendicular to the surface of the globe. The color-coded arcs between the countries show the inter-country traffic, with the higher and redder arcs indicating the larger traffic flows. The globe is illuminated by a light which is positioned to indicate, via the angle of the sun, the time for the frame of the time-series data that is displayed.

The difficulty with general 3D network displays is that they are often confusing and difficult to navigate around, and cause the user to lose a sense of overall context. Restricting the display to a sphere captures many of the advantages of a general 3D network layout while simultaneously helping the user to maintain context. Users are accustomed to globes, so navigation is simplified since the user may rotate the globe interactively and there is little chance of the user becoming disoriented. The number of line crossings, and hence the amount of visual confusion, is also reduced by the three-dimensional embedding and by the presence of the globe surface, which acts as a background.

One weakness of drawing arcs around the globe to show traffic is that only the front side of the globe is visible, making it impossible to see where some arcs terminate. We attempt to overcome this occlusion in three ways. The first uses an interactive translucency control that makes the globe partially transparent. Translucency works well for some examples and some viewing

orientations, but may still be visually confusing if there are a large number of arcs curving around the globe terminating at many different locations. The second allows the user to select from several different ways of routing the arc paths, including routes that run through the middle of the globe. Appropriate selection of the path can reduce the display clutter, and the routing of paths through the globe is effective when used in combination with the translucency control. Finally, the user may perform filtering of the arcs to select only those with certain attributes. Incorporating user interface controls such as filtering and translucency can further reduce the visual complexity of the display, and thereby lead to greater insights.

2.2 Arc Maps

Another related technique draws arcs among nodes positioned on a flat 2D map (see Figure 2) embedded in the 3D space. Each node is positioned geographically, as on the globe, and then arcs are drawn between the nodes with the height, color, and thickness of arc encoding the statistic. As with the global display, drawing the arcs in 3D eliminates the line crossings that curse the 2D displays.

Arc maps have several desirable characteristics. Firstly, unlike the global network display discussed before, the geographic context of the network is not restricted to world-wide networks. Arc maps can display a network on any map at any desired resolution. The retention of geographic context makes the displays interpretable. Secondly, although the geographic map is drawn on a 2D plane, it can be arbitrarily positioned and oriented in the space. Therefore the user can interactively navigate the display by translating, rotating, scaling, and visualizing the network from different perspectives under different rendering conditions. Thirdly, as with the global display, drawing arcs greatly eliminates the line crossings associated with planar 2D displays. Finally, the most important links, such as those having the largest traffic flows, are represented by the highest arcs and therefore are visually predominant from different angles.

By parameterizing the height using a slider, as we do in our implementation, a user may interactively and smoothly transit between a traditional 2D node and link map and a 3D arc map. To further reduce the visual cluttering of the display, we have also incorporated another two options into the arc map. Firstly, the glyphs representing the country encode the total packet

count for all links emanating from the country using both color and size, drawing visual attention to the most important countries. Secondly, the arcs can be rendered with translucency, thereby avoiding any occlusion of otherwise hidden arc segments. Figure 3 presents the same scenario as Figure 2 with these two options.

2.3 Drill-Down Network Views

The visual network display techniques illustrated in Figures 1 and 2 attempt to show a complete network, with all links and nodes visible simultaneously. For many analysis tasks, however, users need to *drill down* and obtain more details about a particular node or subnetwork. For example, in Figure 1 the largest glyph, colored orange and corresponding to the United States, is clearly interesting. Where do the traffic flows from the United States go? Which country has the highest community of interest with the United States? The lowest? From Figure 1 it is difficult to answer fine-grained detail-oriented questions about a particular node, even after rotating and adjusting the viewing parameters.

Linked views showing “details on demand” are common in the data mining community and well established in the statistical graphics community. This subsection describes three types of network-oriented drill-down display for displaying all links emanating from a designated *focal node*.

The first view, shown in Figure 4, looks somewhat like the spokes on a wheel. This spoke view shows all traffic between the USA (the central node) and each other country. The spokes are color-coded and the nodes size-coded to represent the link traffic. In this particular example the nodes are positioned alphabetically in a circle, but the nodes could be ordered according to geography or even placed in geographic position, for example on a polar projection map centered on the focal node. The latter view could also be considered as a version of the arc map in which the graph has been filtered to show only the focal node and the links and nodes directly connected to it.

Spoke displays work well for a small number of nodes (less than 50 or 100) but eventually become overwhelmed as the number of nodes increases and the glyphs begin to overplot. The reason is that spoke displays do not make effective use of screen real-estate because all lines are the same length. We can partially overcome this problem using a 3D layout. The helix view generalizes the spoke view by positioning the nodes on a 3D helix

(see Figure 5). When viewed from above, the helix view is a spoke view. By interactively rotating the helix view all nodes come into view sequentially.

Another 3D layout, motivated by the helix display, positions the nodes approximately uniformly around a sphere as with a pincushion (see Figure 6). In this view the nodes are actually arranged along circles of "latitude", with the number of circles and number of nodes per circle chosen so the angle between circles and between nodes on a circle are approximately the same. An alternate placement algorithm tessellates the surface of the sphere and selects points from the tessellation. To be effective, the pincushion view (like the helix view) needs to be viewed interactively with motion.

3 Discussion

The figures in this paper are screen images created from the SeeNet3D network visualization system. This system, developed over the last few years, is a comprehensive network visualization designed for visual exploration of large, time-varying communication networks. SeeNet3D is the most recent member of a series of network visualization and exploration environments that we have developed in an on-going research program aimed at visual network data analysis. Earlier members include SeeNet [BEW95] and NicheWorks [EW93].

SeeNet displays time-oriented geographic network data using a suite of 2D displays and overcomes the display clutter problem using *dynamic parameter adjustment*. The analyst manipulates the display parameters interactively while watching the display change; good parameter focusing is achieved when the display shows meaningful information about the data. The contribution of SeeNet involves the careful choice of orthogonal dynamic controls and the appreciation that interactivity and dynamic manipulation may usefully address display clutter.

The NicheWorks system is most useful for abstract networks where there is no natural geographic layout. It uses a two-pass "spring-based" algorithm with several heuristics that the user may fine tune interactively. The essential idea in the positioning algorithm is that all nodes repel with unit force and the relationships among them form a counter-balancing attraction. Numerically solving the motion equations results in an incremental positioning that places related nodes close to each other [FR91]. For sparse networks, placing related nodes together has the very desirable effect of reducing the

display clutter caused by long connecting links. NicheWorks embeds the placement algorithms in an interactive data manipulation, control, and exploration environment that exploits many techniques from dynamic statistical graphics [BCW87].

The new contribution of SeeNet3D involves the use of 3D geographic node positioning. Ongoing technology trends in processor speeds and display hardware are making the creation of interesting and useful 3D network displays on desk-top equipment feasible. Such 3D network displays, in our opinion, are more engaging and visually interesting than the 2D displays that were previously standard. There is also accumulating evidence that 3D network displays are more effective than 2D displays [ABW93] [WF96]. SeeNet3D attempts to exploit the engaging aspects of 3D displays while simultaneously using 3D layouts to overcome the fundamental problems inherent in displaying large networks.

The global network view (Figure 1) captures the inherent geography in world-wide networks and displays network traffic using a natural and understandable metaphor. It scores well on the engaging criteria and has been used by other researchers for such purposes as visualizing MBone Internet traffic [MHCF96].

The arc view (Figures 2 and 3) captures many of the benefits of the globe while retaining the advantages of the traditional 2D network displays. Users are able to find the most pleasing viewing angles by interactively manipulating the display. This view may also be easily incorporated into other 3D network display systems, and is not restricted to world-wide networks.

The linked drill-down views (Figures 4, 5, and 6) show all traffic from a designated node. In our implementation the user mouses on an interesting node in either the global or arc view to set the focal node in the drill-down views. The spoke view is conventional and works well for nodes with degree up to perhaps 50 or 100. For nodes with more connections the spokes become too close. The helix and spherical views place the nodes in 3D and use rotation and human perception to address the limited screen space problem and effectively support much larger fanouts than the spoke view. In our experience both the helix and sphere layouts are more effective than the spoke layout, although we prefer the helix layout perhaps because it is more ordered.

4 Implementation

The SeeNet3D system is currently a 5,000 line C++ program built on top of the *Vz* framework. *Vz* is a visualization platform embodied in an object-oriented, cross-platform (MS Windows, OpenGL, and X11) C++ library. The *Vz* library provides a foundation for building highly-interactive, linked-view graphical displays. The *Vz* C++ Library:

- hides platform and operating system differences;
- handles display rendering in a portable manner;
- provides a standard “look and feel”;
- facilitates the view linking;
- includes many utility classes for data management, statistics, and mathematics.

As the foundation for data visualization, the library provides the core and common functions in our system and tools.

We have version of SeeNet3D running on SGI workstations and PCs running MS Windows 95 and Windows NT. On a top-end personal computer (150MZ processor) with graphics accelerator we can render the display in less than a second. This is not yet quite fast enough for interactive performance, which we find requires at least four to five image refreshes per second.

5 Summary

This paper describes five views for showing complex geographic networks embodied in the SeeNet3D system. SeeNet3D is a network visualization environment for the exploration of large, time-varying communication networks. The research focus is on how to overcome some fundamental problems in understanding large and complex networks using 3D graphics technologies, while simultaneously maintaining the benefits of the useful and well-established 2D node and link maps. Two of the views display the complete network. The other three show subnetworks and may be used to “drill down” for details of all links emanating from a designated node.

By applying 3D technologies, we have successfully provided more visual information in a more understandable and engaging manner. Visually, our 3D views avoid some major problems of 2D views and present the graph in a clearer and cleaner way. Equally important, the users have more capability to understand the network through interactive manipulation of the display.

Information Visualization is an important field with great potential. In this paper we have only addressed a few of the issues associated with the 3D display of geographical networks. Many fundamental problems remain to be solved before the whole picture of 3D Information Visualization can be drawn.

Acknowledgments

We gratefully acknowledge and appreciate Audris Mockus's help in providing some of the maps used in our views.

References

- [ABW93] K. W. Arthur, K. S. Booth, and C. Ware. Evaluating 3D task performance for fish tank virtual worlds. *ACM Transactions on Information Systems*, 11(3):239–265, 1993.
- [BCW87] Richard A. Becker, William S. Cleveland, and Allan R. Wilks. Dynamic graphics for data analysis. *Statistical Science*, 2:355–395, 1987.
- [Ber81] Jacques Bertin. *Graphics and Graphic Information Processing*. Walter de Gruter & Co., Berlin, 1981.
- [BETT94] G. Di Battista, P. Eades, R. Tamassia, and I. G. Tollis. Algorithms for drawing graphs: An annotated bibliography. *Computational Geometry Theory and Applications*, 4:235–282, 1994.
- [BEW95] Richard A. Becker, Stephen G. Eick, and Allan R. Wilks. Visualizing network data. *IEEE Transactions on Visualization and Graphics*, 1(1):16–28, 1995.

- [CE95] Kenneth C. Cox and Stephen G. Eick. 3D displays of internet traffic. In Nahum Gershon and Stephen G. Eick, editors, *Proceedings Information Visualization Symposium*, pages 129–131, Atlanta, Georgia, 30 October 1995.
- [CT96] I. F. Cruz and J. P. Twarog. 3D graph drawing with simulated annealing. In F. J. Bredenbrueg, editor, *GD'95 Proceedings Lecture Notes in Computer Science*, volume 1027, pages 162–165. Springer-Verlag, 1996.
- [EW93] Stephen G. Eick and Graham J. Wills. Navigating large networks with hierarchies. In *Visualization '93 Conference Proceedings*, pages 204–210, San Jose, California, 25-29 October 1993.
- [FPF88] Kim M. Fairchild, Steven E. Poltrack, and George W. Furnas. Semnet: Three-dimensional graphic representations of large knowledge bases. In R. Guindon, editor, *Cognitive Science and Its Applications for Human Computer Interaction*, pages 201–233. Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1988.
- [FR91] T. Fruchterman and E. Reingold. Graph drawing by force-directed placement. *Software—Practice and Experience*, 21(11):1129–1164, 1991.
- [GE95] Nahum Gershon and Stephen G. Eick. *Proceedings Information Visualization '95*. IEEE Computer Science Press, 1995.
- [GE96] Nahum Gershon and Stephen G. Eick. *Proceedings Information Visualization '96*. IEEE Computer Science Press, 1996. To appear.
- [GKNV93] Emden R. Gansner, Eleftherios E. Koutsofios, Stephen C. North, and K.P. Vo. A technique for drawing directed graphs. *IEEE Transactions on Software Engineering*, 19(3):214–230, March 1993.
- [Kau91] Arie Kaufman. *Volume Visualization*. IEEE Computer Society Press, Los Alamitos, California, 1991.

- [LR94] John Lamping and Ramana Rao. Laying out and visualizing large trees using a hpyerbolic space. In *UIST: Proceedings of the ACM Symposium on User Interface Software and Technology*, pages 13–14, 1994.
- [MHCF96] Tamara Munzner, Eric Hoffman, K. Claffy, and Bill Fenner. Visualizing the global topology of the mbone. In Nahum Gershon and Stephen G. Eick, editors, *Information Visualization '96*. IEEE Computer Science Press, 1996. To appear.
- [Noi93] E. G. Noik. Layout-independent fisheye views of nested graphs. In *Proceedings of the 1993 IEEE Symposium on Visual Languages*, pages 336–341, 1993.
- [PvW93] F. H. Post and T. van Walsum. Fluid flow visualization. In *Focus on Scientific Visualization*, pages 1–40. Springer, 1993.
- [RCM93] George G. Robertson, Stuard K. Card, and Jock D. Mackinlay. Information visualization using 3D interactive animation. *Journal of the ACM*, 36(4):56–71, 1993.
- [SA94] Lindsey Spratt and Allan Ambler. Using 3D tubes to solve the intersection line representation problem. In *IEEE Symposium on Visual Languages*, pages 254–261, 1994.
- [SB94] Manojit Sakar and Marc H. Brown. Graphical fisheye views. *Communications of the ACM*, 37(12):73–84, December 1994.
- [Tuf83] Edward R. Tufte. *The Visual Display of Quantitative Information*. Graphics Press, Cheshire, Connecticut, 1983.
- [Tuf90] Edward R. Tufte. *Envisioning Information*. Graphics Press, Cheshire, Connecticut, 1990.
- [WF96] Colin Ware and Glenn Frank. Evaluating stereo and motion cues for visualizing information nets in three dimensions. *ACM Transactions on Graphics*, 15(2):121–140, April 1996.

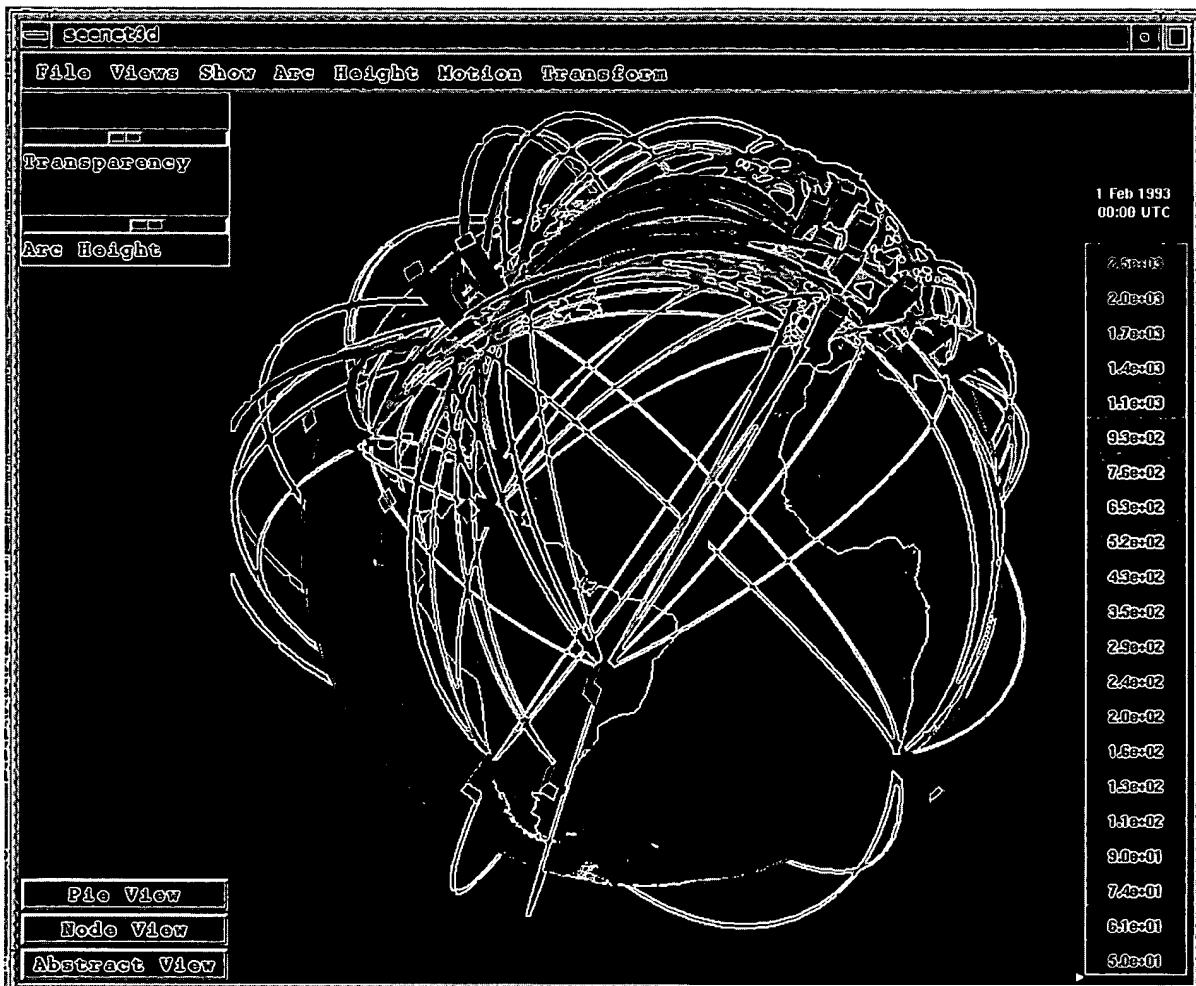


Figure 1: World-wide Internet traffic over a two-hour period, with the color and thickness of the lines encoding the traffic.



Figure 2: An arc map showing the world-wide Internet traffic in Figure 1.



NOT AVAILABLE COPY

Figure 3: Partially translucent arcs avoid the overplotting in Figure 2.

BEST AVAILABLE COPY



Figure 4: Traffic to and from the USA shown as spokes in a wheel. The size and color of the links encodes the link traffic.

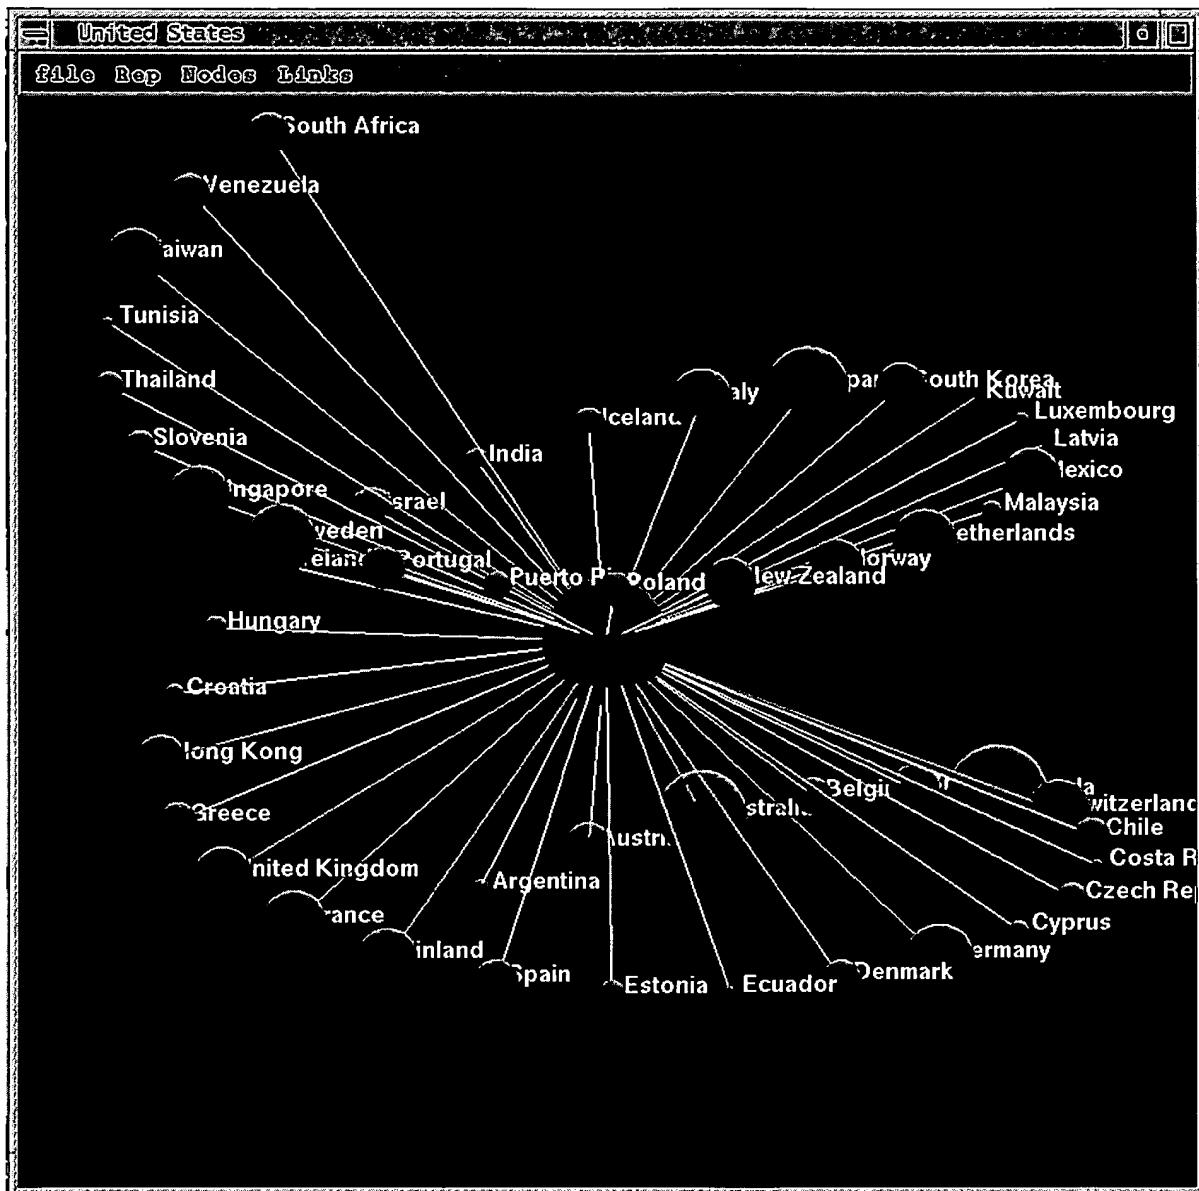


Figure 5: Positioning the nodes in a helix makes better use of screen real-estate than does the spoke layout in Figure 4.

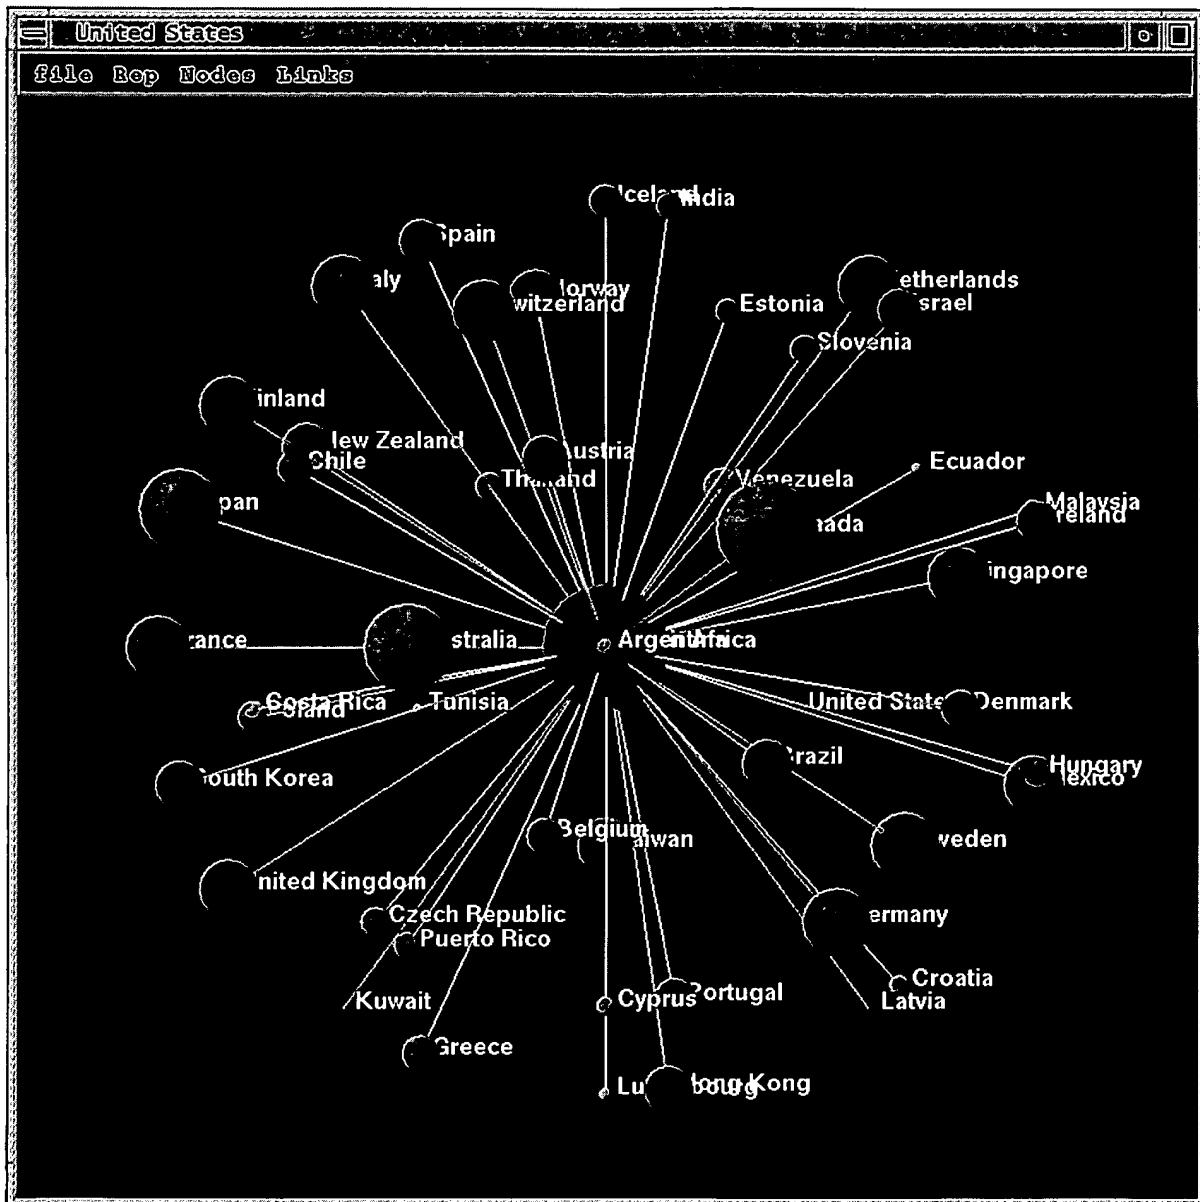


Figure 6: The pincushion view positions the nodes from Figure 4 approximately uniformly on the surface of a sphere.

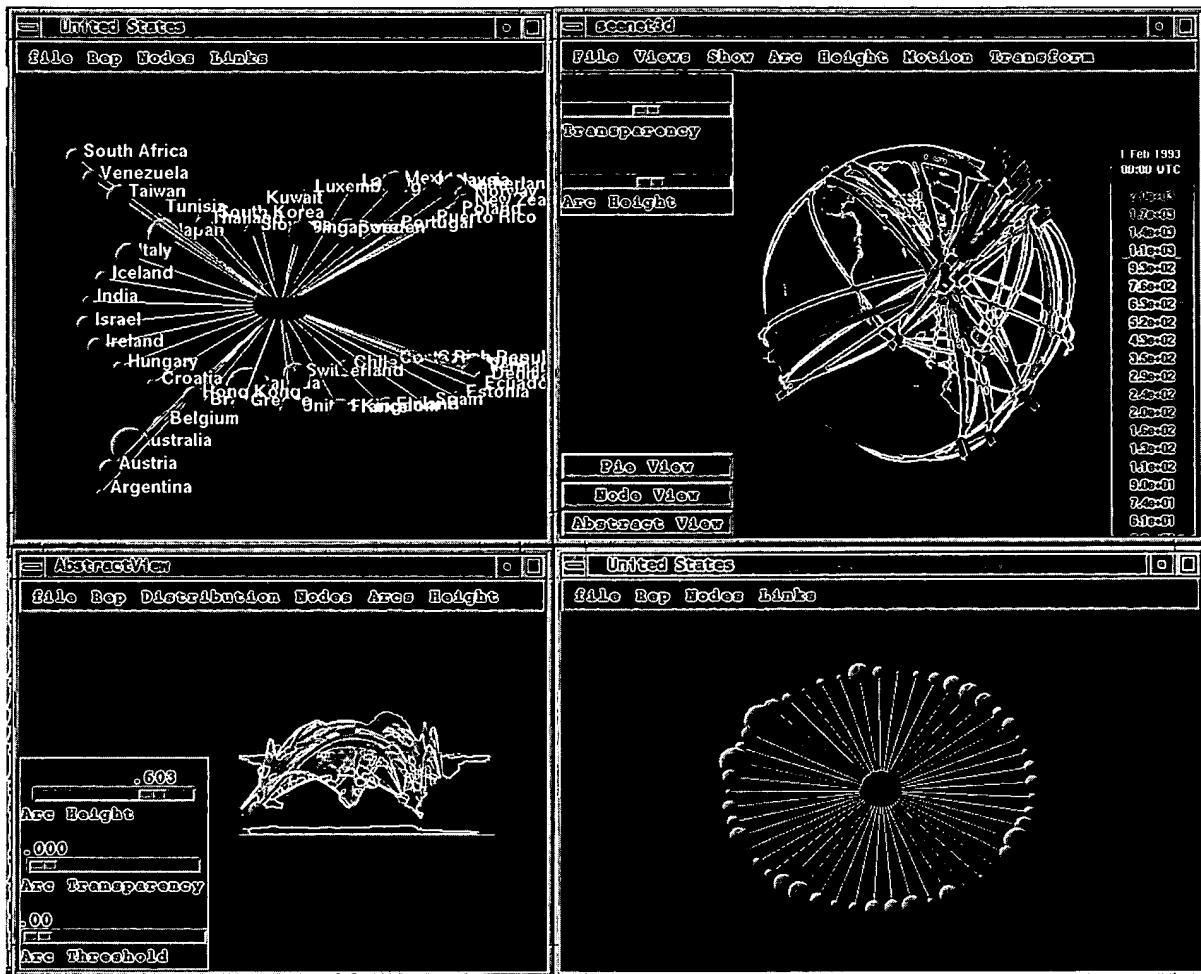


Figure 7: The SeeNet3D network visualization system.